

AD-A103 319 RAND CORP SANTA MONICA CA
AI FOR SYSTEMS MANAGEMENT, (U)
JAN 81 F HAYES-ROTH
UNCLASSIFIED RAND/P-6573

F/G 9/2

NL

1 OF 1
40 A
010321-9

END
DATE
PRINTED
10-81
DTIC

AD A103319

LEVEL 1

1

AI FOR SYSTEMS MANAGEMENT


Frederick Hayes-Roth

SENT
ELT
1/10/81

1/10/81

January 1981

FILE COPY



IPR:ND/P-6573

81 8 25 105

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional Staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

**The Rand Corporation
Santa Monica, California 90406**

1

ACKNOWLEDGMENTS


The work described in this paper has been performed by many persons. The paper describes research and development conducted primarily within the Information Processing Systems research program at The Rand Corporation. In addition to myself, this research has been performed principally by Dan Gorlin, Barbara Hayes-Roth, Phil Klahr, Stan Rosenschein, Don Waterman, and Bob Wesson. The work on speech understanding systems reported herein was conducted at Carnegie-Mellon University and was performed primarily by Lee Erman, Vic Lesser, Bruce Lowerre, Raj Reddy, and myself.

DTIC
FICTE
AUG 26 1981
D
H

| | |
|--|--|
| DISTRIBUTION STATEMENT A | |
| Approved for public release; Distribution Unlimited | |

SUMMARY

Complex systems require intelligent control strategies, and AI concepts and tools may contribute to the management of such systems. At Rand, we have been developing AI approaches to systems management problems. Our work involves three principal components: (1) a model of the system to be managed; (2) a situation assessment function that employs the model to interpret sensor data; and (3) a planning and control function that employs the model to select desired actions. This broad approach generalizes many of the recent advanced AI applications and defines a substantial R&D program. Our current R&D efforts aim at improving the technologies for modeling and simulation, for systematizing and improving situation assessment methods, and for expanding our repertoire of planning strategies and tools. This paper describes these efforts in overview.



CONTENTS

Section

| | |
|---|-----|
| ACKNOWLEDGMENTS | iii |
| SUMMARY | v |
| I. INTRODUCTION | 1 |
| II. SITUATION ASSESSMENT AND PLANNING APPLICATIONS | 4 |
| III. PROGRESS IN MODELING AND SIMULATION | 8 |
| IV. SITUATION ASSESSMENT AS A PROBLEM OF INTERPRETATION | 14 |
| V. MODEL-BASED DYNAMIC PLANNING AND CONTROL | 22 |
| VI. CONCLUSION | 26 |
| VII. REFERENCES | 29 |

I. INTRODUCTION

AI FOR SYSTEMS MANAGEMENT




Figure 1

To manage systems, we need to understand what they are doing and how they will respond to potential interventions. In addition, we need a method for generating and selecting interventions that will produce desirable outcomes. We use the term "situation assessment" to refer to the task of understanding the system's current behavior. We use the term "future projection" to refer to the task of anticipating the likely response of the system to potential interventions. We refer to the process of designing, evaluating, and implementing intervening actions as "planning and control."

We believe that system management is an ideal problem for artificial intelligence. Previous applications of AI have led us to believe that the three management functions discussed above require

knowledge-based, heuristic methods. In particular, we see a primary role for a system model that mediates situation assessment and planning and control.

A system model, in our framework, represents the elements of the domain and their interrelationships. Such a model supports situation assessment by providing a basis for linking observable data to apparent system states through causal or "diagnostic" reasoning. When we observe some data, we attempt to find the possible states of the system which the model suggests would produce those data. In a similar way, a system model can support future projection. In this case, we employ the model as a simulation. Planning requires that we examine one or more alternative simulations to identify a desirable future among our options.

Recent advances in knowledge representation bring us close to the point where we can use one form of system model for all these purposes. In conventional approaches, system models can either support simulation or analysis, but not both. The lack of flexibility in traditional models arises from the particular forms of representation and computing previously available. For example, typical simulations comprise event or time-stepped procedural descriptions of state changes. To answer questions with such models, we specify initial states, simulate all state changes over some period, and then retrospectively analyze the simulator's outputs. In contrast to such traditional approaches, we are moving toward declarative descriptions of modeled entities and their behaviors. These descriptions can be applied for a variety of purposes, including time-stepped or event-stepped simulation, deductive question-

answering, situation assessment, and, when indicated, debugging and model revision.

Much of our current work focuses on improved modeling methods and more systematic techniques for interpretation and forecasting. In the remainder of this paper, I will describe representative projects, accomplishments, and current objectives in these areas.

II. SITUATION ASSESSMENT AND PLANNING APPLICATIONS

Our research group has developed several prototype applications of the model-based approach to situation assessment and planning. These are shown in Table 1.

TABLE 1.

SOME CURRENT APPLICATIONS OF SITUATION ASSESSMENT AND PLANNING

| TASK | MODEL | SITUATION ASSESSMENT | PLANNING |
|-------------------------|--|---|---|
| 1. Air traffic control | Civilian aviation flightplans | Monitoring current traffic and projected collisions | Heuristic flight planning and collision avoidance |
| 2. Tactical Navy | TECA--own and opposing forces (naval platforms) | Monitoring opposing force positions and movements to detect threats | Threat avoidance countermeasures |
| 3. Strategic Air Forces | ROSS--own and opposing air forces and air defenses | Current status of forces and projected battle outcomes | Dynamic force management |
| 4. Civil justice | ROSIE rules for product liability law and cases | Analyzing the basis of current decisions and settlements | Designing and assessing potential legislation |

These four applications are representative of the wide range of systems which may prove fruitful targets for more sophisticated management techniques. In air traffic control, for example, Rand personnel have designed and tested heuristic planning techniques for en route flight control. This prototype system contains rules about how airplanes move and how their flightplans can be altered by controller commands. These rules are used to project ahead and spot potential collisions as well as to evaluate the desirability of heuristically proposed fixes. In the Navy application, the TECA system contains rules about how our own and opposing ships behave and the circumstances under which one force threatens another. This knowledge can be used to look ahead for potential threats as well as exploring the space of alternative actions that avoid these threats. In the strategic application, the ROSS system models our own and opposing air forces to enable faster-than-real-time evaluation of air engagements. This should permit military planners to explore new concepts of dynamic force management. The fourth application above has begun only recently. It aims at modeling part of the civil justice system concerned with product liability lawsuits and settlements. The premises, inferences, and decisions reached in a representative sample of precedent cases define a model of the decisionmaking process. This model is captured in the form of rules written in the ROSIE programming language. This model allows analysts to foresee the likely effects of a change in the law by extrapolating from its inferred effects over a sample of cases. An illustration of a legal rule represented in the ROSIE programming language is shown below:

[RULES: ORDINARY CARE DEFINITION]

If the victim is an adult who does know the proper use of the product
and (the victim does know 'the product is defective'
and the victim does continue the use of the product)
or ((the victim does know 'the product is dangerous'
or the victim does know 'the product is defective')
and (the victim is careless in the use of the product
or the victim is inattentive in the use of the product))
or (the victim is improper in the use of the product
and ((there is a warning by the manufacturer
and that warning does describe the improper use of the product)
or (there is a warning by the seller
and that warning does describe the improper use of the product)))
or (the victim does know 'the victim is sensitive to the product'
and the victim does continue the use of the product)
or the victim does use poor practices in the use of the product
or (there is a means for protection from the hazard of the product
and the victim does not use that means),
assert the use of (the product) by the victim does not involve
ordinary-care.

Each of these applications is implemented as a contemporary AI system. However, each is rudimentary and only a pre-prototype of what an operational system would require. The experimental systems reveal both the strengths and weaknesses of AI tools for managing complex systems. Most of these systems require extensive human involvement. On the other hand, the requirements for planning and control generally

exceed human computing capabilities. As a consequence, we look for salutary combinations of human and machine intelligence. The machine, for example, can project future flight positions of aircraft more accurately and quickly than humans. On the other hand, humans can often prescribe plan revisions that achieve desired results simply and efficiently. Thus, in a domain like air traffic control, we are aiming to articulate the elements of the problem-solving task, to develop AI functions to accomplish subtasks, and to integrate humans and machines in effective, cooperative problem-solving teams. Similar aims guide our work in the other areas cited.

III. PROGRESS IN MODELING AND SIMULATION

AI IN MODELING AND SIMULATION

NEW CAPABILITIES

SYMBOLIC DESCRIPTIONS AND DEDUCTIONS
INTEGRATION OF PHYSICAL PROCESSES CONSTRAINTS
AND BEHAVIORAL MODELS

NEW LANGUAGES

RULE-BASED SYSTEMS: EMYCIN, RITA, ROSIE
IF CONDITIONS THEN ACTIONS
INTELLIGIBILITY, MODIFIABILITY, EXPLANATION
END-USER INVOLVEMENT

NEW ARCHITECTURES

MULTI-LEVEL MODELS FOR ABSTRACTION AND ADAPTIVE
DETAIL COMPUTATION
GOAL-ORIENTED FOCUSING, SAMPLING, AND REPORTING

Figure 2

Recent work in AI has concentrated heavily on improved methods for representing knowledge, building knowledge bases, and encoding heuristic rules. These reflect, in a larger sense, a new approach to modeling. Most previous computer science work in modeling has been tightly coupled to simulation. It is natural therefore that improved techniques for modeling and simulation should go hand in hand.

AI contributes new capabilities, new languages, and new architectures for models. I will review these contributions each in turn.

The new capabilities include non-procedural descriptions and new methods for manipulating modeled entities. The non-procedural, or

"symbolic" descriptions, provide two benefits. First, they allow us to describe many types of entities and behaviors that don't readily admit to mathematical or state-change representations. Second, they support deductive inference and analysis, problem-solving methods of recognized importance. Within the framework of symbolic descriptions, we can construct models of important systems which incorporate mathematical as well as non-mathematical relationships. This provides a basis for integrating models of physical processes, human behavior, and various constraints.

In the area of modeling languages, a host of new tools are emerging. We have focused our efforts on a variety of rule-based systems. Many of these were initially inspired by the MYCIN program for infectious disease diagnosis (Shortliffe, 1976). Recently, a version of MYCIN without the medical knowledge has been developed which is called EMYCIN (Van Melle, 1980). Along similar lines, Rand developed the RITA system (Anderson, 1976). The newest descendant of this tradition is the ROSIE language, a rule-oriented system for implementing expertise. This language provides a general purpose programming system, an interactive programming environment, symbolic modeling and deductive capabilities, and a friendly and easily learned English syntax.

These new languages have been designed with several goals in mind. First, they should simplify the problem of encoding knowledge. Equally important, however, they should provide an intelligible encoding of that knowledge which is accessible to domain-experts and problem-solving practitioners. Intelligible models are a prerequisite of end-user involvement in the modeling process. Because model construction is a

continuing process, these new languages offer the promise of significantly improved and enhanced models deriving from greater user involvement in the modeling process. Moreover, we believe that intelligible computer code, specifically that encoded in English, provides a natural basis for explaining the system's actions.

New architectures for simulations are also arising. In our work, in particular, we face the problem of making complex system models comprehensible to the agents that employ them (people or machines). Thus, although we require extensive detail in our model for some purposes, we need to simplify and abstract the model for many other purposes. This leads naturally to the idea of multi-level models in which the same relationships are represented simultaneously at different levels of precision and aggregation. Such models should also support intelligent use of resources in time-stressed situations. By choosing to model events at the least detailed level possible for a given purpose or at the most detailed level possible given limited computing time, we adapt our computation to an appropriate level of detail.

Another major change in the architecture of simulations concerns the fundamental purpose of simulating a model, namely to answer some question. Ordinarily, a large simulation must execute completely before any answers can be culled from its results. We have turned simulations inside-out, in a sense. We view them as question-answerers that are executed solely to collect needed data. Thus, beginning with a question, we use simulation-management functions to focus the simulator towards the computations required for the question-at-hand. In the strategic air forces model for example, we may choose to ignore whole

geographic regions or kinds of air defenses in answering particular questions.

Within this goal-oriented framework of simulation, we have also introduced the concepts of concurrent, statistical sampling to produce more timely and efficient analyses of the simulated events. The simulator collects only those data relevant to the current questions. Thus, at the moment the simulator completes the computations in focus, the desired results can be produced.

BOMBER PENETRATION EVENTS




Figure 3

```

1326      ILLOC (114, 77, 111111)
1327      IF (T1EST0 .GT. 1) GO TO 200
1328      100 = 1
1329      101 = 0
1330      CALL S1NP1 (LOCNBS + 1, SP1STP, 100, 101)
1331      IF (T1EST0 .GT. 7) GO TO 210
1332      IF (T1EST0 .GT. 1) GO TO 220
1333      210 CALL S1P1TP (T1, LOCNBS (100)), 1W00 (LOCNBS + 1), SP1STP
1334      GO TO 200
1335
1336      C BASE HAS FIGHTERS OF THE REQUIRED TYPE
1337      C CHECK FEASIBILITY
1338
1339      C BUILD BUDGET FIGHTER STATE VECTOR
1340
1341      200 T0 = 0
1342      111 = -1
1343      1TAKH = 0
1344      CALL S1NP1 (SP1STP, 41, LOCBS + 1, LOCBNH + 1, 101)
1345      IF (T1 .GT. 10) 31 GO TO 210
1346      P1ELFT = 1000000
1347      CALL S1NP1 (LOCNBS + 1, CAPLOC, 1W00 (LOCNBS + 1)), T1LOC (T1, T2 + 01),
1348      1TAKH, 1P1STP, 1PLAC, P1ELFT)
1349      IF (T1PLS .GT. 1) 32 GO TO 380
1350      CALL S1P1TP (T1, LOCNBS (100)), 1W00 (LOCNBS + 1), SP1STP
1351      IF (T1AKH = 2) 200, 300, 300
1352
1353      300 CONTINUE
1354      IF (T1LOC (LOCAP + 10) .GT. 0.01)
1355      1000 (LOCAP + 10) = 0.01
1356      1000 (LOCAP + 10) = 1PLAC (T1) - T1W00
1357
1358      C SEND FEASIBLE FIGHTERS
1359
1360      LOCPI = LOC
1361      LOC = LOC + ILLOC (114, 113)
1362      NEED = 1W00 (LOCCAP + 93)
1363      CALL S1NP1 (SP1STP, NEED, LOCBS + 1, LOCPI + 1, 101)
1364      IF (T1 .GT. 10) 31 GO TO 300
1365      1W00 (LOCAP + 93) = 1W00 (LOCAP + 93) - NEED
1366      1W00 (LOCAP + 93) = 1W00 (LOCAP + 93) + NEED
1367      1W00 (LOCPI + 10) = 111100 + 100 (1W00 (LOCPI + 1), T1LOC (T1, T2 + 01),
1368      1TAKH, 1P1STP (T1), 100)
1369      CALL 1W00 (1PLAC, 1W00 (LOCPI + 93), 5)
1370      CALL 1W00 (1W00 (LOCPI + 100) + 370), 1W00 (LOCPI + 370), 43
1371      1W00 (LOCPI + 100) = P1ELFT
1372      P1ELFT = 1W00 (LOCPI + 300)
1373      1W00 (LOCPI + 300) = NEED
1374
1375      C ADJUST FIGHTER ON BASE COUNTERS
1376
1377      100 = 0

```

Figure 4

BEHAVIOR RULES OF OBJECTS

GCI

IF PENETRATOR IS IN MY RADAR RANGE
THEN TELL MY FILTER CENTER ABOUT PENETRATOR
ADD PENETRATOR TO MY LIST OF OBJECTS IN RANGE
MONITOR PENETRATOR THROUGH RADAR

FILTER CENTER

IF PENETRATOR P IS IN RANGE OF GCI G.
THEN ADD P TO MY LIST OF TRACKED PENS
ADD FACT THAT G IS TRACKING P
IF P IS NEW, THEN REQUEST FIGHTER ASSIGNMENT

FIGHTER.

IF PENETRATOR IS IN MY RANGE
THEN DO A MONTE CARLO FOR DETECTION
IF PENETRATOR IS DETECTED THEN FIRE MISSILE

Figure 5

Figures 3, 4, and 5 illustrate some of these new modeling concepts in the strategic air force problem. Figure 3 illustrates the variety of objects, behaviors, and events in a hypothetical U.S.-U.S.S.R. conflict. Figure 4 illustrates some of the Fortran code in an actual 350,000 line program used to model these entities for military planning purposes. Unfortunately, such code cannot be understood by anyone and, as a consequence, it cannot be continually modified to reflect evolving reality. Worse yet, it cannot be scrutinized and contemplated by military experts; thus, it should not be greatly trusted.

In Figure 5, we have illustrated the kind of simulator language our ROSS project is developing. Simplified behavior rules for ground control intercept radars (GCI), command-control filter centers, and air defense fighter interceptors are shown. Each rule, presumably, provides a clear, intelligible model of behavior. Rules quite like these constitute our current model and associated simulation.

IV. SITUATION ASSESSMENT AS A PROBLEM OF INTERPRETATION

In AI, the problem of perception has come to be viewed as a problem of interpreting observations vis-a-vis a model. When this paradigm is applied to systems management, it suggests that to understand what a system is doing we need to interpret its behavior vis-a-vis a model of that system.

Perhaps the best examples of AI systems that perform interpretation tasks come from the domain of speech understanding. While the speech understanding problem is interesting in its own right, in this paper I will introduce it simply to illustrate two different AI approaches to interpretation. Each of these suggests a general approach to interpretation which will find application in different situation assessment contexts. Readers interested in the speech understanding problem should consult Lea (1980) or Erman et al. (1980).

TWO ALTERNATIVE AI INTERPRETATION SYSTEMS: **1. THE HEARSAY-II SPEECH UNDERSTANDING SYSTEM**




Figure 6

TWO ALTERNATIVE AI INTERPRETATION SYSTEMS:
2. THE HARPY SPEECH UNDERSTANDING SYSTEM




Figure 7

I will describe two different systems which understand connected English sentences drawn from a syntactically constrained grammar with a 1000-word vocabulary. These two systems were called Hearsay-II (Erman *et al.*, 1980) and Harpy (Lowerre & Reddy, 1980). These systems differ primarily in the way they organize and control the search for likely interpretations.

To understand speech, a system needs a variety of capabilities. It must transduce the physical signal into an acoustic measurement, usually a waveform relating amplitude to time. Intervals of these acoustic

measurements correspond to spoken sounds, such as phonemes, syllables and words. The speech understanding system must associate the physical parameter measurements with sequences of words that might have produced the observations. In doing this, it must rule out implausible word sequences. The central problem for such systems is considering plausible words and word sequences, evaluating them, and selecting the most probable.

The Hearsay-II system is illustrated in Figure 6. This figure depicts the "blackboard" database structure on which hypothetical interpretations are recorded. The blackboard is a two-dimensional structure which locates hypotheses based on their time of occurrence in the spoken utterance and their interpretation at various levels of abstraction. These levels define a hierarchy of increasingly aggregated interpretations. The lowest level consists of acoustic segments, intervals of speech which manifest relatively unchanging physical measurements. At the next higher level, one or more successive segments may be interpreted as a particular acoustic phone. Our speech systems employed approximately 80 distinct phones. Successive levels aggregate and interpret contiguous hypothetical interpretations from the adjacent lower level. In Hearsay-II, hypotheses could be formed at the lexical level (any of more than 1000 different English words), the word-tuple level (sequences of words that appear grammatical according to a first-order Markov approximation to the grammar), the syntactic level (grammatical phrases), and the semantic level (meaningful sentences).

Figure 6 portrays the collection of hypotheses generated midway through the analysis of "Are there any new reports on space complexity?"

At this point in the analysis, the highest-level, correct hypothesis is the three-word sequence, "...new reports on...". During each interval Hearsay-II usually maintained on the blackboard nearly ten times as many incorrect hypotheses as correct ones. Nevertheless, it considered overall a nearly negligible fraction of all possible hypotheses. Such efficient search is the crux of speech understanding.

In forming an overall interpretation of a sentence, Hearsay-II acted somewhat like an organization of cooperating specialists. Distinct programs performed analyses at the various levels of interpretation. Each specialist looked for hypotheses on the blackboard that could trigger its own inferential capacities to suggest new or modified hypotheses at nearby levels. A managerial specialist watched the overall flow of activity and regulated resource allocations to coordinate and focus the team activity. In this way, Hearsay-II integrated multiple, diverse sources of knowledge in a cooperating problem-solving system.

The overall approach to problem-solving that Hearsay-II employed has been called "opportunistic" (cf. Hayes-Roth & Hayes-Roth, 1979; Nii & Feigenbaum, 1978). Many people believe that such a system mirrors the cognitive processes of human problem-solving both within an individual and within groups of individuals. In fact, the blackboard mechanism was motivated in large part by the desire to allow several human speech experts to work independently in the construction of specialist programs. By restricting the specialists' interactions to blackboard hypothesization, each specialist was relieved of responsibilities for understanding the internal models employed by the others.

The Harpy system, on the other hand, employed a radically different approach to organizing and controlling the search for plausible interpretations. Harpy integrates all levels of speech knowledge into a homogeneous finite state transition network representing the possible sequences of spoken phones. To create such a network, all speech knowledge must be represented either in terms of finite state grammars or corresponding finite state transition graphs. Thus, there is one graph for each word which re-expresses the word in terms of the possible sequences of phones that it can manifest. Similarly, the grammar of English is represented as a network consisting of all possible sequences of words. These two types of networks are combined by successively replacing tokens at one level by their network definitions at the adjacent lower level. This continues until the lowest level representation is reached at which point the entire language is modeled by a huge network of possible lowest-level transitions.

Figure 7 schematizes a small portion of such a network. Along the uppermost path through the network, the successive phones denote one particular pronunciation of "Are there any new reports ...". The figure also depicts the search process Harpy employs. To interpret a speech input, Harpy moves through the network from start to finish following a few of the most plausible paths of interpretation. At each point, it compares the next segment of speech with the possible continuations of its current path interpretations. Each is evaluated for its acceptability relative to the model's expectation. An overall goodness-of-fit measure is associated with the extended path. A small number, usually around 200, of the most probable paths are retained and

the process iterates.

When Harpy reaches the end of the input utterance, it accepts the most probable path it has generated as the interpretation of the sentence. This search method, which has been called "beam search," operates in a manner akin to dynamic programming with pruning. By using stored networks and by pruning unlikely paths from consideration, the beam search technique achieves amazing efficiency.

Figure 8 summarizes the contrasting attributes of Hearsay-II and Harpy. The numbers in Figure 8 aggregate approximately actual performance statistics.

COMPARING HEARSAY-II AND HARPY

| FEATURE | HEARSAY-II | HARPY |
|--------------------------------------|------------------------------------|-------------------------------|
| CONNECTED SPEECH IN 1000 WORDS | YES | YES |
| PROCESSING TIME PER SECOND OF SPEECH | 100 | 15 |
| SENTENCES CORRECTLY UNDERSTOOD | 90% | 95% |
| KNOWLEDGE REPRESENTATION | INDEPENDENT SPECIALISTS | COMPILED FINITE-STATE NETWORK |
| HEURISTIC SEARCH METHOD | OPPORTUNISTIC HYPOTHESIZE-AND-TEST | BEAM SEARCH |

Figure 8

Two important lessons about interpretation tasks should be drawn from this comparative review. First, interpretation lends itself to a variety of methods which vary chiefly in the way they organize, apply, and control knowledge. Second, well-defined problem-solving tasks can often be performed with great efficiency if flexible control strategies are replaced by rigid, systematic control algorithms. Of course, in management science both of these lessons are familiar ones. It seems pleasing that our AI systems may be reaching the level of complexity where the same principles begin to apply.

I have surveyed these speech understanding applications because they provide the most detailed and concrete examples of existing interpretation systems. They also reveal the general ideas that govern situation assessment. Situation assessment requires two key ingredients, a model and a method for assigning meaning to sensor data. In the speech systems, we saw the same underlying knowledge represented in two different ways. In Hearsay-II, the knowledge was segmented into individual procedures that acted within their own levels of representation and analysis. In Harpy, all knowledge was represented in a homogeneous, large transition network. Each of these two representations was engineered to support the kind of control strategy the designers envisioned. In the case of Hearsay-II, the control concept was opportunistic and cooperative problem-solving. In Harpy, it was beam search.

In the future, I anticipate that we will see other types of control strategies with corresponding specialized representations for the system models which underlie the interpretation process.

V. MODEL-BASED DYNAMIC PLANNING AND CONTROL

Our approach to planning and control for managing systems rests heavily on the use of a system model for future projection and option evaluation. Basically, we view planning and control as a continual process of situation assessment and dynamic replanning. The overall process consists of two phases, as shown in Figures 9a and 9b.

AN AI APPROACH TO DYNAMIC REPLANNING




Figure 9a

AN AI APPROACH TO DYNAMIC REPLANNING




Figure 9b

The first phase of planning and control requires us to update our initial model to accommodate sensor data and then to project the likely consequences of currently planned actions. Most models have parameters or alternatives which must be fit to the sensor data. In speech, for example, the models permit 1000 alternative words that may occur during each interval in time. In the case of air traffic control, on the other hand, the parameters stand for the possible identities, intentions, trajectories and current states of each aircraft to be controlled. Our models enable us to predict what our sensors should see under various

alternatives, and we find plausible those interpretations that are consistent with the data. Situation assessment tunes our model to fit the current data.

At this point, we use the model to look-ahead into the future. Assuming all entities behave according to the plans and processes prescribed by the model, we can anticipate the likely consequences under current assumptions. In the case of air traffic control, for example, we can look to see if any likely collisions will occur. In general, we simulate specifically to search for projected threats and opportunities.

The next figure closes the replanning loop in an obvious way. We take the projected threats and opportunities as inputs to a replanning process. Such a process generates hypothetical new plans. For example, we might imagine turning both planes in a projected collision ninety degrees from their current bearing. Such hypothetical plans give rise to revised, hypothetical models. These in turn may again be simulated to project potentially undesirable consequences. In turn, we may need to refine further the hypothetical plan. Or, worse yet, the plan may prove so unpromising that we may abandon it and explore alternative kinds of fixes.

For such dynamic replanning to work we need two capabilities. The first is a method for generating alternative plans. The second is a fast method for projecting the hypothetical plan's consequences. Toward the first objective, we are developing a variety of planning mechanisms with particular heuristic methods in various domains (e.g., air traffic control, tactical air targeting, route planning). Toward the second objective, we are developing new architectures for rapid, goal-oriented

simulation, as previously described.

It seems likely that systematic methods for modeling, situation assessment, and dynamic replanning will emerge during the next decade of AI research.

VI. CONCLUSION

CURRENT STATUS AND FUTURE PROSPECTS

MODELING

VERY HIGH-LEVEL LANGUAGES
HYBRID PHYSICAL & BEHAVIORAL MODELING
MODELS AS KNOWLEDGE BASES

SITUATION ASSESSMENT

SPEECH AND IMAGE UNDERSTANDING
 SYSTEM MONITORING
 INTELLIGENCE ANALYSES
 GENERAL PURPOSE PROCEDURES

PLANNING

OPPORTUNISTIC INITIAL PLAN FORMULATION
DYNAMIC MONITORING AND PROJECTION
SYSTEMATIC, ITERATIVE PLAN REFINEMENT

Figure 10

Figure 10 summarizes many of the points covered in this paper. It also suggests some of the possible benefits that we can expect from this line of work. For example, we anticipate that as modeling languages grow in sophistication and intelligibility, models will become valuable assets that systematize important bodies of knowledge.

I have surveyed a broad area of research in this very short paper. This raises the significant danger of oversimplification and glibness. In fact, the field of AI has just begun to scratch the surface of systems management. However, the early results seem quite promising as well as generalizable. They also reinforce the belief that success in

AI applications requires a blend of theory, domain specific knowledge, and engineering. In the absence of any apparent obstacles, we look forward to continued steady progress on all fronts.

VII. REFERENCES

Anderson, Robert H. and Gillogly, J. J. Rand intelligent terminal agent (RITA): design philosophy. R-1809-ARPA, The Rand Corporation, Santa Monica, California, 1976.

Erman, L. D., Hayes-Roth, F., Lesser, V.R., and Reddy, D. R. The Hearsay-II speech understanding system: integrating knowledge to resolve uncertainty. Computing Surveys, 1980,

Gorlin, D., Hayes-Roth, F., Rosenschein, S., Sowizral, H., Waterman, D.
A. Programming in ROSIE: an introduction by means of examples.
N-1646-ARPA. The Rand Corporation, Santa Monica, California,
February 1981.

Gorlin, D., Hayes-Roth, F., Rosenschein, S., Sowizral, H., Waterman, D.
A. Rationale and Motivation for ROSIE. N-1648-ARPA. The Rand Corporation, Santa Monica, California, February 1981.

Gorlin, D., Hayes-Roth, F., Rosenschein, S., Sowizral, H., Waterman, D.
A. The ROSIE language reference manual, N-1647-ARPA. The Rand Corporation, Santa Monica, California, February 1981.

Hayes-Roth, B. and Hayes-Roth, F. A cognitive model of planning.
Cognitive Science, 1979, 3, 275-310.

Klahr, Philip and Faught, W. S. Knowledge-Based Simulation. In Proceedings of the First Annual National Conference on Artificial Intelligence, Stanford University, 1980.

Lea, Wayne A. (ed.), Trends in Speech Recognition, Prentice-Hall, Englewood Cliffs, New Jersey, 1980.

Lowerre, Bruce and Reddy, D. R. The Harpy speech understanding system. In Wayne Lea (ed.), Trends in Speech Recognition. Prentice-Hall, Englewood Cliffs, New Jersey.

Nii, H. P. and Feigenbaum, E. A. Rule based understanding of signals. In Donald Waterman and Frederick Hayes-Roth (eds.), Pattern-Directed Inference Systems, Academic Press, New Jersey, 1978.

Shortliffe, E. H., Computer-based Medical Consultations: MYCIN, American Elsevier, New York, 1976.

Waterman, D. A., Anderson, R. H., Hayes-Roth, F., Klahr, P., Martins, G., Rosenschein, S. J., Design of a rule-oriented system for implementing expertise. N-1158-1-ARPA. The Rand Corporation, Santa Monica, California, May 1979.

Waterman, D. A. and Peterson M. Models of Legal Decisionmaking. R-2717-ICJ, The Rand Corporation, Santa Monica, California, 1981.

Waterman, D. A. and Peterson, M. Rule-based models of legal expertise. In Proceedings of the First Annual National Conference on Artificial Intelligence, Stanford University, 1980.

Wesson, Robert. Planning in the world of the air traffic controller. Proceedings of the 5th International Joint Conference on Artificial Intelligence, MIT AI-Lab, Cambridge, Massachusetts, 1977.

van Melle, W. A domain-independent production-rule system for
consultation programs. Proceedings of the 6th International
Joint Conference on Artificial Intelligence, August 1979.

